

**IN THE UNITED STATES PATENT AND TRADEMARK OFFICE**

Applicant: Baccelli et al.

Examiner: Bokhari, Syed M.

Serial No.: 10/721,399

Art Unit: 2616

Filed: November 25, 2003

Docket No.: YOR920030277US1 (8728-634)

For: **APPARATUS FOR SCALABLE RELIABLE GROUP COMMUNICATION**

**Mail Stop AF**

Commissioner for Patents

P.O. Box 1450

Alexandria, VA 22313-145

**RESPONSE UNDER 37 CFR 1.116**

Examiner:

This reply is in response to the Final Office Action dated October 29, 2008. Please consider the following remarks.

**CLAIMS**

## Listing of claims

1. (Previously Presented) A computer-implemented method for group communication over a network of processors comprising:

determining an overlay spanning tree comprising an origin node and at least one receiving node;

determining a configuration of the overlay spanning tree having a maximum throughput among all possible configurations of the overlay spanning tree, wherein determining the configuration of the overlay spanning tree comprises

defining a target bandwidth less than a maximum link bandwidth of edges of the overlay spanning tree given a fully connected overlay distribution graph,

constructing a reduced overlay distribution graph by iteratively removing an edge from a current overlay distribution graph, beginning with the fully connected overlay distribution graph, the edge having a bandwidth less than or equal to the target bandwidth, increasing the target bandwidth upon determining that the configuration of the overlay spanning tree is constructible based on the current overlay distribution graph, and decreasing the target bandwidth upon determining that the configuration of the overlay spanning tree is not constructible based on the current overlay distribution graph, until the configuration of the overlay spanning tree has the maximum throughput with no edge having a bandwidth below the target bandwidth; and

controlling a source communication rate between the origin node and the at least one

receiving node to be less than or equal to a bottleneck rate of the configuration of the overlay spanning tree based on the reduced overlay distribution graph.

2. (Original) The computer-implemented method of claim 1, further comprising protecting data delivery by link error recovery.

3. (Original) The computer-implemented method of claim 2, wherein the overlay spanning tree comprises a plurality of nodes, wherein the data delivery is reliable such that each node receives the same data.

4. (Original) The computer-implemented method of claim 1, further comprising scaling the overlay spanning tree to an arbitrary group size.

5-6. (Canceled)

7. (Original) The computer-implemented method of claim 1, further comprising joining a new node to the spanning tree.

8. (Original) The computer-implemented method of claim 7, comprising joining the new

node to an existing node of the spanning tree upon determining that the existing node has a bandwidth of greater than or equal to an existing rate.

9. (Original) The computer-implemented method of claim 8, further comprising:

determining a triangular improvement upon determining that no existing node has a bandwidth greater than or equal to the existing rate;

joining the new node at an attachment point having a highest bandwidth among existing nodes of the spanning tree upon determining that the triangular improvement failed; and  
redetermining the spanning tree upon determining bandwidth less than or equal to a minimum threshold.

10. (Original) The computer-implemented method of claim 1, further comprising  
redetermining the spanning tree upon determining that an existing node has left the spanning tree.

11. (Original) The computer-implemented method of claim 10, further comprising:

determining orphaned child nodes of the existing node that has left the spanning tree; and  
performing a join for each orphaned child node.

12. (Previously Presented) A program storage device readable by machine, tangibly embodying a program of instructions executable by the machine to perform method steps for group communication over a network of processors, the method steps comprising:

determining an overlay spanning tree comprising an origin node and at least one receiving node;

determining a configuration of the overlay spanning tree having a maximum throughput among all possible configurations of the overlay spanning tree, wherein determining the configuration of the overlay spanning tree comprises

defining a target bandwidth less than a maximum link bandwidth of edges of ~~for~~ the overlay spanning tree given a fully connected overlay distribution graph,

constructing a reduced overlay distribution graph by iteratively removing an edge from a current overlay distribution graph, beginning with the fully connected overlay distribution graph, the edge having a bandwidth less than or equal to the target bandwidth, increasing the target bandwidth upon determining that the configuration of the overlay spanning tree is constructible based on the current overlay distribution graph, and decreasing the target bandwidth upon determining that the configuration of the overlay spanning tree is not constructible based on the current overlay distribution graph, until the configuration of the overlay spanning tree has the maximum throughput with no edge having a bandwidth below the target bandwidth; and

controlling a source communication rate between the origin node and the at least one receiving node to be less than or equal to a bottleneck rate of the configuration of the overlay spanning tree based on the reduced overlay distribution graph.

13. (Original) The method of claim 12, further comprising protecting data delivery by link error recovery.

14. (Original) The method of claim 13, wherein the overlay spanning tree comprises a plurality of nodes, wherein the data delivery is reliable such that each node receives the same data.

15. (Original) The method of claim 12, further comprising scaling the overlay spanning tree to an arbitrary group size.

16-17. (Canceled)

18. (Original) The method of claim 12, further comprising joining a new node to the spanning tree.

19. (Original) The method of claim 18, comprising joining the new node to an existing node of the spanning tree upon determining that the existing node has a bandwidth of greater than or equal to an existing rate.

20. (Original) The method of claim 19, further comprising:  
determining a triangular improvement upon determining that no existing node has a bandwidth greater than or equal to the existing rate;  
joining the new node at an attachment point having a highest bandwidth among existing nodes of the spanning tree upon determining that the triangular improvement failed; and  
redetermining the spanning tree upon determining bandwidth less than or equal to a minimum threshold.

21. (Original) The method of claim 12, further comprising redetermining the spanning tree upon determining that an existing node has left the spanning tree.

22. (Original) The method of claim 21, further comprising:  
determining orphaned child nodes of the existing node that has left the spanning tree; and  
performing a join for each orphaned child node.

23. (Canceled)



**REMARKS**

Claims 1-4, 7-15 and 18-22 are pending in the present application. Applicants respectfully request reconsideration of the application in view of the above remarks.

**I. Rejections Under 35 U.S.C. § 103**

Claims 1-2, 4, 10, 12-13, 15 and 21 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *McCanne* (US 2003/0088696 A1) in view of *Haas et al.* (US 7,035,937 B2).

Claims 1 and 12 are the independent claims.

Claims 1 and 12 claim, *inter alia*, “defining a target bandwidth less than a maximum link bandwidth of edges of the overlay spanning tree given a fully connected overlay distribution graph.”

Applicants appreciate the Response to Arguments found in the Final Office Action. The Examiner stated essentially that the “bandwidth constraints specified by external policies” is analogous to the limitation above. In response, please consider the following:

*McCanne* teaches methods for controlling packet flow through points of bandwidth constraint specified by external policies (see paragraph [0044]). *McCanne* does not teach or suggest a method for configuring an overlay spanning tree, much less, defining a target bandwidth less than a maximum link bandwidth of edges of the overlay spanning tree given a fully connected overlay distribution graph, as essentially claimed in Claims 1 and 12. The bandwidth constraints of *McCanne* are policy-defined bandwidth constraints (see paragraph [0047]) that form choke points in a network – *McCanne*’s intended purpose is to transform (thin) a stream of packets as necessary to pass through the choke points. Thus, *McCanne*’s bandwidth constraint is

not analogous to the claimed target bandwidth; the claimed target bandwidth is used as a threshold to remove edges having bandwidths below the threshold bandwidth (choke points) in a process of configuring an overlay spanning tree, while *McCanne*'s policy-defined bandwidth constraints are choke points in a network through which the packet flow is made to fit. Thus, *McCanne* changes the packet data to fit the bandwidth constraints of an existing network while the claimed invention creates a network (through edge removal) to support at least the threshold bandwidth – clearly then these are quite different inventions. For example, in response to a choke point, *McCanne* transforms the packet flow by reducing frame rates or dropping packets (see paragraphs [0010 and 0044]), whereas the claimed invention would remove the choke point from the overlay spanning tree (assuming it is less than the threshold bandwidth).

Furthermore, the reduction in path length cited in the rejection is a routing choice and is unrelated to the claimed invention. Consider that the reduction in path length is the result of the addition of connections between LVIFs, such that the LVIFs may peer with each other (see paragraph [0065]). Consider further that a path length may be increased by removing an edge as claimed – the claimed invention is directed to creating a network having at least a target bandwidth, not necessarily a shortest path.

For at least the foregoing reasons, *McCanne* does not teach or suggest every limitation of Claims 1 and 12.

*Haas* teaches a routing protocol for an ad hoc network that employs alternate tree computation algorithms that continually compute backup trees that can be employed to replace failed trees. *Haas* does not teach or suggest “defining a target bandwidth less than a maximum link bandwidth of edges of the overlay spanning tree given a fully connected overlay distribution

graph,” as claimed in Claims 1 and 12. *Haas* is silent on implementing bandwidth requirements or constraints of any kind. Therefore, *Haas* fails to cure the deficiencies of *McCanne*.

The combination of *McCanne* and *Haas* teaches bandwidth constraints specified by external policies and a routing protocol for an ad hoc network that employs alternate tree computation algorithms that continually compute backup trees that can be employed to replace failed trees. The combination does not teach or suggest “defining a target bandwidth less than a maximum link bandwidth of edges of the overlay spanning tree given a fully connected overlay distribution graph,” as claimed in Claims 1 and 12. Accordingly, the combination does not teach or suggest every limitation of Claims 1 and 12.

Claims 1 and 12 further claim, *inter alia*, “constructing a reduced overlay distribution graph by iteratively removing an edge from a current overlay distribution graph, beginning with the fully connected overlay distribution graph, the edge having a bandwidth less than or equal to the target bandwidth.”

The Examiner point to paragraph 65, lines 1-13 of *McCanne* as teaching constructing a reduced overlay distribution graph by removing an edge from the fully connected overlay distribution graph having a bandwidth less than or equal to the target bandwidth. Applicants respectfully disagree. Respectfully, *McCanne* teaches reducing the path lengths between a leaf node and an external network to improve routing performance and using other paths as backups in case designated links go down and (see paragraph 65, lines 3-13). The reduction of path lengths of *McCanne* merely refers to choosing a path having fewer links among a plurality of available paths, while the remaining paths are used for backup and peering. As shown above, the reduction

in path length is the result of additional connections in the form of peer relationships between the LVIFs. Nowhere does *McCanne* teach or suggest removing an edge having a bandwidth less than or equal to a target bandwidth; all links forming the paths of *McCanne* are considered in routing decisions. Selecting a path for its path length and using other paths as backups is clearly not analogous to removing edges from an overlay distribution graph with a bandwidth less than a target bandwidth. Therefore, *McCanne* fails to teach or suggest every limitation of Claims 1 and 12.

*Haas* teaches a routing protocol for an ad hoc network that employs alternate tree computation algorithms that continually compute backup trees that can be employed to replace failed trees. *Haas* does not teach or suggest iteratively removing an edge from a current overlay distribution graph, beginning with the fully connected overlay distribution graph, the edge having a bandwidth less than or equal to the target bandwidth, as essentially claimed in Claims 1 and 12. Nowhere does *Haas* discuss a target bandwidth, let alone iteratively removing edges having a bandwidth less than or equal to a target bandwidth. Therefore, *Haas* fails to cure the deficiencies of *McCanne*.

The combination of *McCanne* and *Haas* teaches choosing a single edge for use with an LVIF when multiple edges are incident to the LVIF and using the remaining edges for backup and peering, and a routing protocol for an ad hoc network that employs alternate tree computation algorithms that continually compute backup trees that can be employed to replace failed trees. The combination does not teach or suggest “constructing a reduced overlay distribution graph by iteratively removing an edge from a current overlay distribution graph, beginning with the fully connected overlay distribution graph, the edge having a bandwidth less than or equal to the target

bandwidth,” as claimed in Claims 1 and 12. Accordingly, the combination does not teach or suggest every limitation of Claims 1 and 12.

Therefore, for at least the above reasons, Claims 1 and 12 are believed to be patentable and non-obvious over the combination of *McCanne* and *Haas*. Applicants respectfully submit that inasmuch as Claims 2, 4, 10, 13, 15 and 21 are dependent on Claims 1 and 12, the dependent claims are allowable for at least the reasons given for Claims 1 and 12. Reconsideration of the instant rejection is respectfully requested.

Claims 3, 11, 14 and 22 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *McCanne* in view of *Haas*, in further view of *Silton et al.* (US 6,327,252).

Claims 3, 11, 14 and 22 depend from Claims 1 and 12, and are believed to be allowable for at least the reasons given for Claims 1 and 12. Reconsideration of the instant rejection is respectfully requested.

Claims 7-9 and 18-20 stand rejected under 35 U.S.C. § 103(a) as being unpatentable over *McCanne* in view of *Haas*, in further view of *Hsu* (US 6,363,319 B1), in further view of *Grover et al.* (US 2002/0187770 A1).

Claims 7-9 and 18-20 depend from Claims 1 and 12, and are believed to be allowable for at least the reasons given for Claims 1 and 12. Claim 23 has been canceled. Reconsideration of the instant rejection is respectfully requested.

**CONCLUSION**

For the forgoing reasons, the application, including Claims 1-4, 7-15 and 18-22, is believed to be in condition for allowance. Early and favorable reconsideration of the case is respectfully requested.

Respectfully submitted,

Date: December 10, 2008

By: /Nathaniel T. Wallace/  
Nathaniel T. Wallace  
Reg. No. 48,909  
Attorney for Applicant(s)

**F. Chau & Associates, LLC**  
130 Woodbury Road  
Woodbury, New York 11797  
TEL: (516) 692-8888  
FAX: (516) 692-8889